Efficient Routing for Cost Effective Scale-out Data Architectures

Ashwin Narayan
Williams College
ashwin.narayan@williams.edu

Vuk Marković
University of Novi Sad
mvukmarko@gmail.com

Natalia Postawa
Adam Mickiewicz University in Poznań
np96924@st.amu.edu.pl

Anna King
University College Cork
111396801@umail.ucc.ie

Alejandro Morales
University of California, Los Angeles
ahmorales@math.ucla.edu

K. Ashwin Kumar
Veritas Labs
ashwin.kayyoor@veritas.com

Petros Efstathopoulos
Symantec Research Labs
petros_efstathopoulos@symantec.com

arXiv:1606.08884v1 [cs.DB] 28 Jun 2016

***Abstract*—Efficient retrieval of information is of key importance when using Big Data systems. In large scale-out data architectures, data are distributed and replicated across several machines. Queries/tasks to such data architectures, are sent to a router which determines the machines containing the requested data. Ideally, to reduce the overall cost of analytics, the smallest set of machines required to satisfy the query should be returned by the router. Mathematically, this can be modeled as the *set cover* problem, which is NP-hard, thus making the routing process a balance between optimality and performance. Even though an efficient greedy approximation algorithm for routing a single query exists, there is currently no better method for processing multiple queries than running the greedy set cover algorithm repeatedly for each query. This method is impractical for Big Data systems and the state-of-the-art techniques route a query to all machines and choose as a cover the machines that respond fastest. In this paper, we propose an efficient technique to speedup the routing of a large number of real-time queries while minimizing the number of machines that each query touches (*query span*). We demonstrate that by analyzing the correlation between known queries and performing query clustering, we can reduce the set cover computation time, thereby significantly speeding up routing of unknown queries. Experiments show that our incremental set cover-based routing is 2.5 times faster and can return on average 50% fewer machines per query when compared to repeated greedy set cover and baseline routing techniques.**

I. Introduction

Large-scale data management and analysis is rapidly gaining importance because of an exponential increase in the data volumes being generated in a wide range of application domains. The deluge of data (popularly called "Big Data") creates many challenges in storage, processing, and querying of such data. These challenges are intensified further by the overwhelming variety in the types of applications and services looking to make use of Big Data. There is growing consensus that a single system cannot cater to the variety of workloads, and different solutions are being researched and developed for different application needs. For example, column-stores are optimized specifically for data warehousing applications, whereas row-stores are better suited for transactional workloads. There are also hybrid systems for applications that require support for both transactional workloads and data analytics. Other systems are being developed to store different types of data, such as document data stores for storing XML or JSON documents, and graph databases for graph-structured or RDF data.

One of the most popular approaches employed in order to handle the increasing volume of data is to use a cluster of commodity machines to parallelize the compute tasks (*scale-out* approach). Scale-out is typically achieved by partitioning the data across multiple machines. Node failures present an important problem for scale-out architectures resulting in data unavailability. In order to tolerate machine failures and to improve data availability, data replication is typically employed. Although large-scale systems deployed over scale-out architectures enable us to efficiently address the challenges related to the volume of data, processing speed and and data variety, we note that these architectures are prone to resource inefficiencies. Also, the issue of minimizing resource consumption in executing large-scale data analysis tasks is not a focus of many data systems that are developed to date. In fact, it is easy to see that many of the design decisions made, especially in scale-out architectures, can typically reduce overall execution times, but can lead to inefficient use of resources [1][2][5]. As the field matures and the demands on computing infrastructure grow, many design decisions need to be re-visited with the goal of minimizing resource consumption. Furthermore, another impetus is provided by the increasing awareness that the energy needs of the computing infrastructure, typically proportional to the resource consumption, are growing rapidly and are responsible for a large fraction of the total cost of providing the computing services.

To minimize the scale-out overhead, it is often useful to control the unnecessary spreading out of compute tasks across multiple machines. Recent works [1][2][5][3][4] have demonstrated that minimizing the number of machines that a query or a compute task touches (*query span*) can achieve multiple benefits, such as: *minimization of communication overheads*, *lessening total resource consumption*, *reducing overall energy footprint* and *minimization of distributed transactions costs*.

A. The Problem

In a scaled-out data model, when a query arrives to the query router, it is forwarded to a subset of machines that contain the data items required to satisfy the query. In such a data setup, a query is represented as the subset of data needed for its execution. As the data are distributed, this implies that queries need to be routed to multiple machines hosting the necessary data. To avoid unnecessary scale-out overheads, the size of the set of machines needed to cover the query should be minimal [1][2][5]. Determining such a minimal set is mathematically stated as the *set cover* problem, which is an NP-hard problem. The most popular approximate solution of the set cover problem is a *greedy algorithm*. However,

running this algorithm on each query can be very expensive or unfeasible when several million queries arrive all at once or in *real-time* (one at a time) at machines with load constraints. Therefore, in order to speed up the routing of queries, we want to reuse previous *set cover* computations across queries without sacrificing optimality. In this work, we consider a generic model where a query can be either a database query, web query, map-reduce job or any other task that touches set of machines to access multiple data items.

There is a large amount of literature available on single query *set cover* problems (discussed in Section II). However, little work has been done on sharing *set cover* computation across multiple queries. As such, our main objective is to design and analyze different algorithms that efficiently perform *set cover* computations on multiple queries. We catalogue any assumptions used and provide guarantees relating to them where possible, with optimality discussions and proofs provided where necessary to support our work. We conducted an extensive literature review looking to understand algorithms, data structures and precomputing possibilities. This included algorithms such as *linear programming* [6], data structures such as *hash tables*, and precomputing possibilities such as *clustering*. We built appropriate models for analysis that included a model with nested queries, and a model with two intersecting queries. These models afforded us a better understanding of the core of the problem, while also allowing us to test the effectiveness of our methods and tools. We developed an algorithm that will solve *set cover* for multiple queries more efficiently than repeating the greedy *set cover* algorithm for each query, and with better optimality than the current algorithm in use (see Section VII-A2). We propose algorithm frameworks that solve the problem and experimentally show that they are both fast and acceptably optimal. Our framework essentially analyzes the history of queries to cluster them and use that information to process the new incoming queries in real-time. Our evaluation of the clustering and processing algorithms for both frameworks show that both sets are fast and have good optimality.

The key contributions of our work are as follows:

- To the best of our knowledge, our work is the first to enable sharing of *set cover* computations across the input sets (queries) in real-time and amortize the routing costs for queries while minimizing the average *query span*.
- We systematically divide the problem into three phases: *clustering the known queries*, *finding their covers*, and, with the information from the second phase, *covering the rest of the queries* as they arrive *in real time*. Using this approach, each of the three phases can then be improved separately, therefore making the problem easier to tackle.
- We propose a novel entropy based real-time clustering algorithm to cluster the queries arriving in real-time to solve the problem at hand. Additionally, we introduce a new variant of greedy *set cover* algorithm that can cover a query Q_i with respect to another correlated query Q_j.
- Extensive experimentation on real-world and synthetic datasets shows that our incremental *set cover*-based routing is $2.5\times$ faster and can return on an average 50% fewer machines per query when compared to repeated greedy *set cover* and baseline routing techniques.

The remainder of the paper is structured as follows. Sections II and III present related work and problem background. In Section IV we describe our query clustering algorithm. Section V explains how we deal with the clusters once they are created and processing of real-time queries in Section VI. Finally, Section VII discusses the experimental evaluation of our techniques on both real-world and synthetic datasets, followed by conclusion.

II. Related Work

SCHISM by Curino et al., [3] is one of the early studies in this area that primarily focuses on minimizing the number of machines a database transaction touches, thereby improving the overall system throughput. In the context of distributed information retrieval Kulkarni et al., [4] show that minimizing the number of document shards per query can reduce the overall query cost. The above related work does not focus on speeding up query routing. Later, Quamar et al., [2][1] presented SWORD, showing that in a scale-out data replicated system minimizing the number of machines accessed per query/job (*query span*) can minimize the overall energy consumption and reduce communication overhead for distributed analytical workloads. In addition, in the context of transactional workloads, they show that minimizing query span can reduce the number of distributed transactions, thereby significantly increasing throughput. In their work, however, the router executes the greedy *set cover* algorithm for each query in order to minimize the *query span*, which can become increasingly inefficient as the number of queries increases. Our work essentially complements all the above discussed efforts, with our primary goal being to improve the query routing performance while retaining the optimality by sharing the *set cover* computations among the queries.

There are numerous variants of the *set cover* problem, such as an online *set cover* problem [7] where algorithms get the input in streaming fashion. Another variant is, k*-set cover* problem [8] where the size of each selected set does not exceed k. Most of the variants deal with a single universe as input [9][10][6], whereas in our work, we deal with multiple inputs (queries in our case). Our work is the first to enable sharing of *set cover* computations across the inputs/queries thereby improving the routing performance significantly.

In this work, in order to maximize the sharing of *set cover* computations across the queries, we take advantage of correlations existing between the queries. Our key approach is to cluster the queries so that queries that are highly similar belong in the same cluster. Queries are considered highly similar if they share many of their data points. By processing each cluster (instead of each query) we are able to reduce the routing computation time. There is rich literature on clustering queries to achieve various objectives. Baeza-Yates et al., [11] perform clustering of search engine queries to recommend topically similar queries for a given future query. In order to analyze user interests and domain-specific vocabulary, Chuang et al., [12] performed hierarchical Web query clustering. There is very little work in using query clustering to speed up query routing, while minimizing the average number of machines per query for scale-out architectures. Our work provides one of the very first solutions in this space.

Another study [13] describes the search engine query clustering by analyzing the user logs where any two queries is said to be similar if they touch similar documents. Our approach follows this model, where a query is represented as a set of data items that it touches, and similarity between queries

is determined by the similar data items they access.

III. Problem Background

Mathematically, the set cover problem can be described as follows: given the finite universe Q, and a collection of sets $\mathscr{M} = \{S_1, S_2, \ldots, S_m\}$, find a sub-collection we call cover of Q, $\mathscr{C} \subseteq \mathscr{M}$, of minimal size, such that $Q \subseteq \bigcup \mathscr{C}$. This problem is proven to be *NP-hard* [10]. Note that a brute force search for a minimal cover requires looking at 2^m possible covers. Thus instead of finding the optimal solution, approximation algorithms are used which trade optimality for efficiency [14][6]. The most popular one uses a greedy approach where at every stage of the algorithm, choose the set that covers most of the so far uncovered part of Q, which is a $\ln n$ approximation that runs in $O(n)$ time.

The main focus of this work is the *incremental set cover problem*. Mathematically, the only difference from the above is that instead of covering only one universe Q, set covering is performed on each universe Q_i from a collection of universes $\mathscr{Q} = \{Q_1, Q_2, \ldots, Q_N\}$. Using the greedy approach separately on each Q_i from $\mathscr{Q}$ is the naïve approach, but when N is large, running the greedy algorithm repeatedly becomes unfeasible. We take advantage of information about previous computations, storing information and using it later to compute remaining covers faster.

In this paper, elements in the universe are called *data*, sets from $\mathscr{M}$ are *machines* and sets from $\mathscr{Q}$ are *queries*. Realistically, it can be assumed that data are distributed randomly on the machines with replication factor of r. In this work, we take advantage of the fact that, in real world, queries are strongly correlated [15][16] and enable sharing *set cover* computations across the queries. Essentially, this amortizes the routing costs across the queries improving the overall performance while reducing the overall cost of analytics by minimizing the average *query span* for the given query workload [1][2].

IV. Query Clustering

In order to speedup the *set cover* based routing, our key idea is to reduce the number of queries needed to process. More specifically, Given N queries, we want to cluster the them into m groups ($m << N$) so that we can calculate set cover for each cluster instead of calculating set cover for each query. Once we calculate set cover for each cluster, next step would be to classify each incoming real-time query to one of the clusters and re-use the pre-computed set cover solutions to speedup overall routing performance. To do so, we employ clustering as the key technique for precomputation of the queries. An ideal clustering algorithm would cluster queries that had large common intersections with each other; it would also be scalable since we are dealing with large numbers of queries. In order to serve real-time queries we need an incoming query to be quickly put into the correct cluster.

Most of the clustering algorithms in the literature require the number of clusters to be given by the user. However, we do not necessarily know the number of clusters beforehand. We also want to be able to theoretically determine bounds for the size of clusters, so our final algorithm can have bounds as well. To that effect, we developed entropy-based real-time clustering algorithm. Using entropy for clustering has precedent in the literature (see [17]). Assume that we have our universe U of size n, let K be a cluster containing queries $Q_1, \ldots, Q_m$. Then we can define the probability p_j of data item j being in the cluster K:

$$p_j(K) = \frac{1}{|K|} \sum_{i=1}^{|K|} \chi_j(Q_i) \tag{1}$$

where the characteristic function χ_j is defined by:

$$\chi_j(Q) = \begin{cases} 1, & j \in Q \\ 0, & j \notin Q \end{cases} \tag{2}$$

Then we can define the *entropy* of the cluster, $S(K)$ as

$$S(K) = -\sum_{j=1}^{n} p_j(K) \log_2 p_j(K) + (1 - p_j(K)) \log_2 (1 - p_j(K)) \tag{3}$$

This entropy function is useful because it peaks at $p = 0.5$ and is 0 at $p = 0$ and $p = 1$. Assume we are considering a query Q and seeing if it should join cluster K. For any data element $j \in Q$, if most of the elements in K do not contain j, then adding Q to K would increase the entropy; conversely if most of the elements contain j, then adding Q would decrease the entropy. Thus, minimizing entropy forces a high degree of similarity between clusters.

A. The `simpleEntropy` Clustering Algorithm:

We developed a simple entropy-based algorithm (pseudocode shown in Algorithm 1). As each query Q comes in, we compute the entropy of placing the query in each of the current clusters and keep track of the cluster which minimizes the *expected entropy*: given clusters $K_1, \ldots K_m$ in a clustering $\mathscr{K}$, the expected entropy is given by:

$$\mathbb{E}(\mathscr{K}) = \frac{1}{m} \sum_{j=1}^{m} |K_j| \cdot S(K_j) \tag{4}$$

If this entropy is above the *threshold* described below, the query starts its own cluster. Otherwise, the query is placed into the cluster that minimizes entropy.

Suppose we are trying to decide if query $Q = \{x_1, \ldots, x_n\}$ should be put into cluster K. Let p_i be the frequency with which x_i is in the clusters of K. Then define the set

$$T(Q, K) = \{x_i \in Q \ : \ p_i \geq \theta_1\}$$

for some threshold θ_1. We say that Q is *eligible* for placement in C if $|T(Q, C)| \geq \theta_2 |Q|$ for some other threshold θ_2. Essentially, we say that Q is eligible for placement in K only if "most of the elements in Q are common in K," where "most" and "common" correspond to θ_1 and θ_2 and are user-defined. Of course, we should have $0 \leq \theta_1, \theta_2 \leq 1$. Then, given a clustering $\mathscr{K}$ with clusters $K_1, \ldots, K_m$, we create a new cluster for a query Q only when Q is not eligible for placement into any of the K_i. This forces most of the values in the query to 'agree' with the general structure of the cluster.

The goal is an algorithm that generates clusters with low entropy. Let us say that a *low-entropy cluster*, a cluster for which more than half the data elements contained in it have probability at least 0.9, is a *tight* cluster. The opposite is a *loose* cluster, i.e. many elements have probability close to 0.5.

B. Analysis of the `simpleEntropy` Clusters:

We take a more in-depth look at the type of clusters that form with an entropy clustering algorithm. The first question considered was whether the algorithm is more likely to generate one large cluster or many smaller clusters. To do

Algorithm 1 A simple entropy based real-time clustering algorithm. Here $\mathscr{Q}$ is the list of queries, and θ_1 are θ_2 are the threshold parameters defined in Section IV

```
function SIMPLEENTROPY(𝒬, θ1, θ2)
    𝒞 ← ∅                          ▷ 𝒞 is the set of clusters
    S ← 0                          ▷ S holds the current expected entropy
    M ← 0                          ▷ M is the number of clustered elements
    for Q ∈ 𝒬 do
        S_Q ← ∞  ▷ S_Q is the min expected entropy due to Q
        C_Q ← ∅
        for C ∈ 𝒞 do
            if NOT ELIGIBLE(Q, C) then
                Continue.
            end if
            if EXPECTEDENTROPY(Q,C) < S_Q then
                S_Q ← EXPECTEDENTROPY(Q,C)
                C_Q ← C
            end if
        end for
        if S_Q < ∞ then
            ADDTOCLUSTER(Q,C_Q)
        else
            NEWCLUSTER(Q)
        end if
    end for
    return 𝒞
end function
```

this, we considered how the size of a cluster affects whether a new query is added to it. We are also interested in how the algorithm weights good and bad fits: If a query Q contains data elements $\{y, x_1, \dots x_n\}$, we want to determine how many of the x_i need to be common in a cluster K to outweigh y being uncommon in K.

The setup is as follows. Assume that we have already sorted m queries into a clustering $\mathscr{K}$, and assume the expected entropy of this clustering is Ω. Given m clusters $\{K_1, \dots, K_m\}$, where each cluster K_j has entropy $E(K_j)$, the expected entropy is given in Equation 4.

Now, we want to calculate the change in expected entropy when a new query, Q, is added to a cluster, as a function of the cluster's size and composition. As a simple start, we only consider the change due to a single data entry. Since the total entropy is additive, understanding the behavior due to one data entry helps understand where the query is allocated.

Let $\Delta\mathbb{E}_i$ be the change in expected entropy due to the presence or absence of data element i. Let p_i be the probability value for element i in a given cluster, and let p_i^* be the new probability if the query is added to the cluster. We have:

$$p_i^* = \begin{cases} \dfrac{np_i + 1}{n+1}, & i \in Q \\ \dfrac{np_i}{n+1}, & i \notin Q \end{cases} \tag{5}$$

Let us also define $S(p_i)$ as the entropy of a single data element, i.e.

$$S(p_i) = -p_i \log_2 p_i - (1-p_i)\log_2(1-p_i) \tag{6}$$

Proposition 1. With the above pre-requisites, we can derive that the difference in expected entropy due to data element i by adding a query to a cluster of size n which had a probability p_i for element i, is:

$$\Delta\mathbb{E}_i(\mathscr{K}) = \frac{1}{M+1}(M\Omega - nS(p_i) + (n+1)S(p_i^*)) - \Omega \tag{7}$$

Proof: The derivation is as follows: To get the new expected entropy, we need to remove the old weighted entropy of the cluster, which is given by $nS(p_i)$ and add back the new weighted entropy of that cluster, given by $(n+1)S(p_i^*)$. And now there are $M+1$ total elements so we divide by $M+1$ to obtain the new expected entropy. Then to get the difference, we simply subtract the old expected entropy, Ω. ■

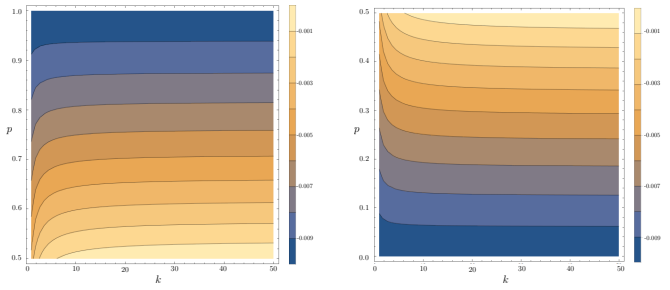

(a) $i \in Q$ (b) $i \notin Q$

Fig. 1. Here we plot Equation 7 setting the expected entropy $\Omega = 1$ and the total number of queries processed $M = 100$, but the general shapes remain true for a wide variety of parameters. The x-axis is the number of queries in the cluster under consideration and the y-axis is the probability of element i in that cluster. We see that the entropy difference is independent of cluster size for large enough clusters.

Figure 1 shows some results of this analysis. As clusters become large enough, $\Delta\mathbb{E}$ (the change in expected entropy), becomes constant with respect to the size of the cluster, and change steadily with respect to p. This is desirable because it means that the change in expected entropy caused by adding Q is dependent only on p_i and not on the size of the cluster.

As will be shown in Section V it is important to our cluster processing scheme that all the queries in a cluster share a large common intersection. That is, there should be data elements that are contained in all queries; and when a new query arrives, that query should contain all those shared elements. This property can be given a practical rationale: the data elements present in the large common intersection are probably highly correlated with each other (i.e. companies that ask for one of the elements probably ask for them all), so it makes sense that an incoming query would also contain all of those data elements if it contained any of them.

From a theoretical perspective, we can see that there is a high entropy penalty when an incoming query does not contain the high probability elements of a cluster.

Proposition 2. Assume we have a cluster K of m data elements, all of which have probability p. A query Q is being tested for fit in K. Let us say that Q contains all but km elements of the cluster K. Then, the entropy of adding Q to cluster K is:

$$\begin{aligned}\Delta\mathbb{E}(\mathscr{K}) = \frac{1}{M+1}(M\Omega - nmE(p) + (n+1)kmE\left(\frac{pn}{n+1}\right) \\ +(n+1)(1-k)mE\left(\frac{pn+1}{n+1}\right)) - \Omega\end{aligned} \tag{8}$$

where, as in (7), $\mathscr{K}$ is the clustering, M is the number of data elements processed, Ω is the previous expected entropy

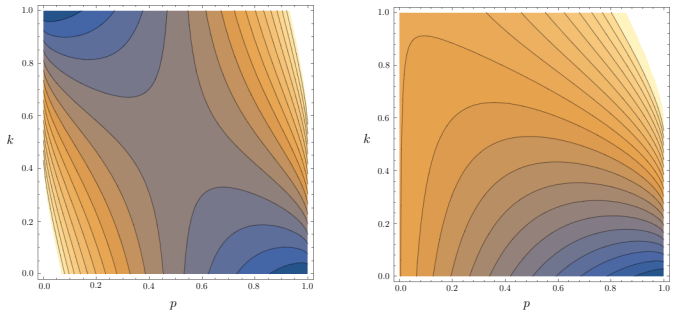

(a) (8) for the full range of probabilities and disjointedness.

(b) (8) concentrating on high probabilities and small k

Fig. 2. A summary of the theoretical basis for the effectiveness of `simpleEntropy`. In (a) we see that all contours become steeper as we reach probability extremes. The central area is a zone of mediocre clustering. In (b) we zoom in on the high probability region. Here we see that most high probability data items *must* be contained in an incoming query.

of the cluster, and n is the number of queries in the cluster.

Proof: The derivation of the above is also analogous to that of (7). The total entropy of the old clustering is given by $M\Omega$, from which we subtract out the old weighted entropy of cluster K, given by $n \cdot mE(p)$, since $E(p)$ is the entropy contributed by each of the m data elements of K, and we weight this by n, the size of the cluster. Then we add the new weighted entropy of K after Q is added to it: There are now km data elements which have probability $pn/(n+1)$ (since Q does not contain these); and there are $(1-k)m$ data elements which have probability $(pn+1)/(n+1)$ (since Q does contain these). Finally, since we are looking for the change in expected entropy, we subtract out Ω, the old expected entropy. ■

In Figure 2, we plot Equation 8 holding n and m at 20 (a typical value for clusters with our parameters according to our experiments.) Figure 2(a) actually provides a general overview of the clustering landscape for an incoming query. We can divide the graph into three regions of interest: (1) the bottom-right and top-left ($k \approx 1-p$), or the *high-quality* region; (2) the central region ($k \approx p \approx 0.5$), or the *mediocre* region and (3) the top-right and bottom-left ($k \approx p \approx 1$ or 0), the *low-quality* region. We discuss the implication of each region in turn:

High-quality Region: These are the points of lowest entropy in the plot. In the bottom-right of Figure 2(a), the implication is that the cluster is a tight cluster to begin with (most elements in it have high probability), and the incoming query also contains most of these elements. The heavy penalty for not containing these high probability elements reinforces the tightness of the cluster. We zoom in on this scenario in Figure 2(b), as it is the most important scenario for tight clusters. Specifically, imagine a cluster has many elements of probability p, where p is close to 1. Call this set of elements X. Now, when a new query Q comes in, if $X \not\subset Q$ there is a high entropy cost to putting Q in the cluster. This means that the high probability regions of a cluster K are likely conserved as new queries come in. As will be seen, this scenario is highly desirable for the real-time case.

The scenario at the top-left of Figure 2(a) is this: the set of data elements that are under consideration are not contained in a given cluster K, and the incoming query also does not contain most of these elements. This case is also important for describing clustering dynamics. A cluster that has been built already may contain some queries that have a few data elements with low probability. Let X be such a set of data elements with p close to 0. Then a new query coming in is *penalized* for containing *any* of those elements. Those elements will remain with low probability in the cluster. The implication is that once a cluster has some low probability elements, these elements will remain with low probability.

Mediocre Region: This is potentially the most problematic region of the entire clustering. Consider a cluster where most of elements have probability p close to 0.5. Then when a new query comes in, the entropy change is similar regardless of how many elements of the cluster the query contains. And if the query is put into this cluster, then on average the probabilities should remain close to 0.5, so the cluster type is perpetuated. Such a cluster is not very tight.

Low-quality Region: This region describes when an ill-fitting query is put into a tight cluster. For example, if we have a cluster for which all the elements have probability p close to 1 and a query comes in which contains none of these elements, we would be in the top-right region of Figure 2(a). This is clearly the worst possible situation in terms of a clustering. Therefore it is a desirable feature of our clustering algorithm that the entropy cost for this scenario is relatively high.

From analyzing these regions, we see that our algorithm will generate high-quality and mediocre clusters (from the high-quality and mediocre regions respectively), but the entropy cost of the low-quality region is too high for many low-quality clusters to form.

V. Cluster Processing

Once our queries are clustered, the goal is to effectively process the clusters as a whole instead of processing each query individually. To that end, we first introduce our so-called `BetterGreedy` algorithm, which is a modified version of the standard greedy algorithm more suited to this problem.

A. *The* `BetterGreedy` *algorithm*

Recall that the standard greedy algorithm covers a query Q with a small number of machines. The `BetterGreedy` algorithm is performed on a query Q_1 *with respect to* another query Q_2. The pseudocode is given in Algorithm 2. At stage k, let $Q_k \subset Q_1$ be the still uncovered elements of Q_1. We choose the machine M^* that contains the most elements of Q_k. In the standard greedy algorithm, if there is a tie, an M^* is chosen arbitrarily. In `BetterGreedy`, if there is a tie, we choose M^* so that it *also* maximizes the elements covered in Q_2. See Figure 3(a) for a visual example.

Algorithm 2 The `BetterGreedy` algorithm. We cover query Q_1 with respect to Q_2. $\mathscr{D}$ is a dictionary, where $\mathscr{D}[i]$ returns the list of machines whose intersection with the uncovered portion of Q_1 is size i, and the list is sorted by intersection size with $Q_2 \setminus Q_1$.

```
function BetterGreedy(Q1, Q2, D)
    M ← ∅                                   ▷ M is the list of machines
    Q ← Q1                                  ▷ Q is the uncovered portion of Q1.
    while |Q| > 0 do
        M ← D.Max   ▷ The machine with the most uncovered elements
        M.Add(M)
        Q ← Q \ M
        D.Update(Q)                         ▷ Update the intersection sizes with Q
    end while
    return M
end function
```

Algorithm Type	Q_1	Q_2	Uncovered Part
Cover just Q_2	1.06+ G_1	G_2	
Cover Q_1 with greedy	G_1	G_2 + .41	1.80
Cover Q_1 with `BetterGreedy`	G_1	G_2 + .05	1.31

TABLE I. A COMPARISON OF THE THREE METHODS FOR PROCESSING THE CLUSTER $Q_1 \subset Q_2$. G_1 GIVES THE SIZE OF THE GREEDY COVER OF Q_1, AND G_2 GIVES THE SIZE OF THE GREEDY COVER OF Q_2. NOTE THAT USING `BetterGreedy` TO PROCESS Q_1 GIVES US A SOLUTION FOR Q_2 THAT IS NEARLY ALWAYS AS GOOD AS THE GREEDY SOLUTION.

B. Analysis of the `BetterGreedy` Algorithm

To make the algorithm as fast as possible, we have a dictionary of lists called `sets_of_size`. Each key in this dictionary is the size of the intersection of each machine with the current uncovered set (and this dictionary is updated at each stage). Corresponding value is a set of all machines of that size. If there are multiple machines under the same key (i.e. they have the same size with respect to the uncovered elements of Q_1), greedy set cover algorithm breaks tie by choosing random machine within a particular size key. However, in the case of our `BetterGreedy` algorithm, they are sorted according to the size of their intersection with $Q_2 \setminus Q_1$. While this additional sorting makes the algorithm worse than standard greedy approach in the worst case (since all the machines could be the same size in Q_1), in practice, our clustering strives to make $Q_2 \setminus Q_1$ small, and so the algorithm is fast enough.

Proposition 3. The above described algorithm runs in $O(\sum_{k=1}^{m} |M_k|) = O(r \cdot |Q|)$, where r is the replication factor when data is distributed evenly on the machines.

Proof: There are two main branches of our algorithm: either under the selected key in $sets_of_size$ is an empty or nonempty set. In the first case, we move our counter to one key below. There are at most $|Q|$ keys in this dictionary, so this branch, which we call a *blank step*, will be performed at most $O(|Q|)$ times, and since it takes $O(1)$, the total time for this branch is $O(|Q|)$. In the other branch, when the set under selected key is not an empty set, we see in the description of the algorithm that in the innermost loop a data unit from a machine is removed (and some other things that all take $O(1)$ are performed). Since there are $\sum_{k=1}^{m} |M_k|$ data units in all machines combined, this part runs in $O(\sum_{k=1}^{m} |M_k|)$. To conclude, the whole algorithm runs in $O(\sum_{k=1}^{m} |M_k| + |Q|) = O(\sum_{k=1}^{m} |M_k|)$. ∎

C. Processing Simple Clusters

In this section we describe the most basic clusters and our ways to process them.

Nested Queries: Consider the most simple query cluster: just two queries, Q_1 and Q_2, such that $Q_1 \subset Q_2$. One might suggest to simply find a cover only for Q_2, using the greedy algorithm, and use it as a cover for both Q_1 and Q_2. In practice, this approach does not perform well. This algorithm is unacceptable in terms of optimality of the cover when comparing the size of the cover for Q_1 given by this approach to the size of the cover produced by the greed algorithm (see Table I). Figure 3(a) with caption explains how our approach solves the problem judiciously.

Intersecting Queries: Here we consider a simple cluster with two queries, Q_1 and Q_2 such that $Q_1 \cap Q_2 \neq \emptyset$. We will do the following: 1) Cover $Q_1 \cap Q_2$ using `BetterGreedy` with respect to $Q_1 \cup Q_2$ 2) Cover the uncovered part of Q_1 and Q_2 separately, using the standard greedy algorithm 3) For the cover of Q_1 return the union of the cover for intersection and uncovered part of Q_1, and for the cover of Q_2 we will give the union of the cover for the intersection and the uncovered part of Q_2. By doing so, we run the `BetterGreedy` once and greedy algorithm twice instead of just running the greedy algorithm twice. However, in the first case, those two greedy algorithms and the `BetterGreedy` algorithm are performed on a smaller total size than two greedy algorithms in the second case. Our algorithm never processes the same data point twice, while the obvious greedy algorithm on Q_1 and Q_2 does. In terms of optimality of the covers obtained in this way, they are on average 0.15 machines (each) larger than the covers we would get using the greedy algorithm. Figure 3(b) gives a visual representation of the algorithm.

D. The General Cluster Processing Algorithm (GCPA)

Using ideas from the previous sections, we developed an algorithm for processing any cluster. We call it the General Cluster Processing Algorithm (`GCPA`). The algorithm, in the simplest terms, goes as follows: 1) Assign a value we call *depth* to each data unit appearing in queries in our cluster. The depth of a data element is the number of queries that data unit is in. For example, consider a visual representation of a cluster on Figure 4(a). On the same figure under Figure 4(b) shows depths of different parts of the cluster. 2) Divide all data units into sets we call *data parts* according to the following rule: two data items are in the same data part if and only if they are contained in exactly the same queries. This will partition the union of the cluster. Also, we keep track of which parts make up each query (which we store in a hash table). 3) We cover the data parts with our desired algorithm (greedy, `BetterGreedy`, ...) 4) For each query we return the union of covers of data parts that make up that query as its cover.

This algorithm can process any shape of a given cluster and allows for a choice of the algorithm used to cover separate data parts. The big advantage of this algorithm is that each data unit (or data part) will be *processed only once*, instead of multiple times as it would be if we were to use the greedy algorithm on each query separately. While dividing the cluster up into its constituent data parts is intensive, this is all pre-computing, and can be done at anytime once the queries are known. Again, it is important to note that our algorithms rely on being able to perform pre-computing for them to be effective.

Since `BetterGreedy` chooses machines that cover as many elements in the cluster as a whole as possible, the covers of the data parts overlap, and makes their union smaller. One thing that can also be used in our favor is that when covering a certain part, we might actually cover some pieces of parts of smaller depths, as illustrated in Figure 4(c). Then, instead of covering the whole of those parts, we can cover just the uncovered elements. Figure 4 shows how this work s step by step. This version of `GCPA`, in which we use the gre edy algorithm, we call `GCPA_G`.

Another option would be to use `BetterGreedy` for processing the parts. The `BetterGreedy` algorithms is done on the data parts with respect to the union of all queries containing that data part and is called `GCPA_BG`. As we will see in the next section, this algorithm gives a major improvement in the optimality of the covers compared to `GCPA_G`.

VI. QUERY PROCESSING IN REAL-TIME

In our handling of real-time processing, we assume that we know everything about a certain fraction of the incoming

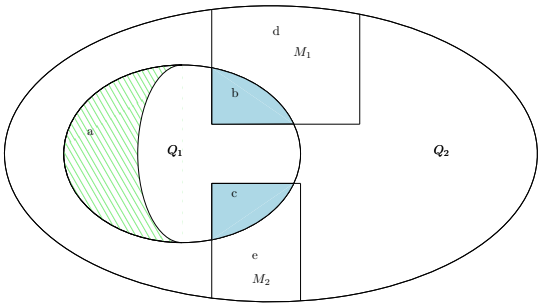

(a) Illustration of `BetterGreedy`

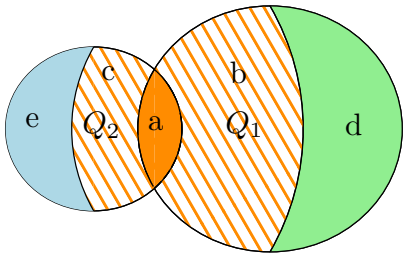

(b) Visual representation of the query intersection algorithm

Fig. 3. In **(a)**, we see an example of `BetterGreedy`. Assume region **a** has already been covered, and we have two machines M_1 and M_2 that cover regions **b** and **c** of Q_1 respectively, and let these regions be the same size. While the standard greedy algorithm would pick from M_1 and M_2 randomly, `BetterGreedy` chooses M_1 because the size of region **d** is bigger than the size of region **e**. The representation of `BetterGreedy` in **(b)** demonstrates that we run the algorithm on the intersection (region **a**) and the striped sections (regions **b** and **c**), thus covering Q_1 and Q_2. Then we simply run the greedy algorithm on the remaining uncovered sections (regions **d** and **e**) to get the full covering. So, for example, the covering of Q_1 is given by the coverings of regions **a**, **b**, and **d**. With standard parameters, we find that covers are on average 0.15 machines larger than the greedy cover.

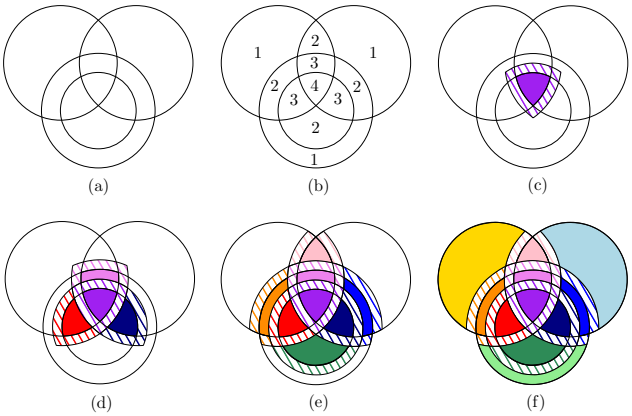

Fig. 4. Example of our algorithm for general cluster processing. In (a) we see the initial state of our cluster of 4 queries. In (b) we see the calculated depths. From (c)-(f) we see in color the part we are processing, and in falling color pattern the cover we end up actually getting. This example shows that instead of doing the greedy algorithm 4 times, we end up doing it 11 times. However, the total size of the data that we are performing our algorithm on is much smaller than doing greedy 4 times because of the overlap in the queries (our algorithm never processes the same part more than once)

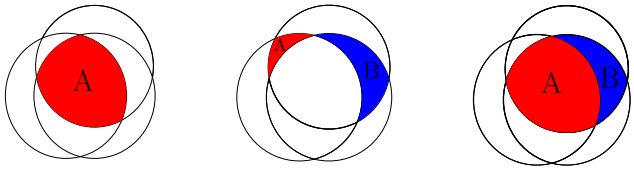

Fig. 5. In the first picture the red area (A) is a part. In the second picture (A) and (B) areas are different parts, because, although they have the same depth, they are made from the intersection of different queries. In the third picture (A) and (B) areas are different parts, as they have different depth.

queries beforehand (call this the *pre-real-time* set), and we get information about the remaining queries only when these queries arrive (the *real-time* set). In other words, we have no information about future queries. To process query in real time the algorithm uses clusters formed in pre-processing.

A. The Real-time Algorithm

Our strategy in approaching this problem is to take advantage of the real-time applicability of the `simpleEntropy` clustering algorithm. From Section VII-B2, we know from experiments, that we only need to process a small fraction of incoming queries to generate most of our clusters. Thus, we cluster the pre-real-time set of queries, and run one of the `GCPA` algorithms on the resulting clusters, storing some extra information which will be explained below. Then, we use this stored information to process the incoming queries quickly with a degree of optimality, as we explain below.

We start by recalling the definition of a *data part* (from Section V-D) and defining the related *G-part*. Given a cluster K and subset of queries in that cluster P, a *data part* is the set of all elements in the intersection of the queries in P but not contained in any of the queries in $K \setminus P$. This implies that all the elements in the same data part have the same depth in the cluster. Figure 5 helps explain the concept.

After running one of the `GCPA` algorithms, the cluster is processed from largest to smallest depth. The *G-part* p_i^g is the set of elements in the cover produced when `GCPA` covers all elements in part p_i that are not in any previous G-part. Note that G-parts also partition our cluster.

To manage processing queries in real time, the algorithm makes use of queries previously covered in clustering process. An array T is created such that for each data item it stores the G-part containing this data item. In other words, element $T[i]$ is the G-part containing element i. For each G-part, we also store the machines that cover that G-part (this information is calculated and stored when `GCPA` is run on the non-real-time queries). The last data structure used in this algorithm is a hash table H, which stores, for each data element, a list of machines covering this data item. In each step the algorithm checks which G-part contains each data item and then for each G-part it checks which machines cover this G-part. Then those machines are added to the set of solutions (if they are not yet in this set). Then for each data item, which was not taken into consideration in any G-part, the algorithm checks in a hash table if any of the machines in the set of solutions covers the chosen data element. In the last step, we cover any still uncovered elements with the greedy algorithm. Data elements on which the greedy algorithm is run form a new G-part.

When a query Q of a length k comes in, the goal is to quickly put Q into its appropriate cluster, so the above algorithm can be run. Since the greedy algorithm is linear in the length of the query, if it takes more than linear time to put a query into a cluster, our algorithm would be slower than

just running the greedy algorithm on the query. Thus, we need to develop a faster method of putting a query into a cluster. We implemented a straightforward solution. Instead of looking at all $O(k)$ potential clusters a query could go into, we just choose one of the elements of the query at random, and choose one of the potential clusters it is in at random. We call this the *fast* clustering method, as opposed to the *full* method (which is $O(k^2)$).

VII. Experimental Evaluation

A. Setup

1) Datasets: We run our experiments on both synthetic and real-world datasets. The sizes considered in this work are the following: each data unit is replicated 3 times and we consider cluster of 50 homogenous machines.

Synthetic Dataset: The total number of data items that we consider is 100K. We generate about 50K queries with certain correlation between them, and each query accesses between 6 and 15 data items. We note that all experiments in this section are done by averaging the results from 1M runs. Following is an explanation of the correlated query generation:

Correlated Query Workload Generation: A sample set of queries is needed to test the effectiveness of a set cover algorithm. As mentioned in Section III the data is distributed randomly on the machines and the queries are correlated. To generate these queries we use *random graphs*. In this context, vertices counted by n represent data and edges represent relations of the data.

A random graph is a graph with n vertices such that each of the $\binom{n}{2}$ possible edges has a probability p of being present. We choose n and p such that $np < 1$, since, in the Erdős-Rényi model [18], this gives a graph with many small connected components. The $np < 1$ case is helpful in this setting because the graph is naturally partitioned into several components as expected. The random graph could, for example, represent a database that contains data from many organizations, so each organization's data is a connected component and is separate from the others. We use a modified DFS algorithm on the random graph to generate nearly highly correlated random queries.

According to the Erdős-Rényi Model for a random graph [18], depending on the value of n and p, there are three main regimes as shown in Figure 6:

- if $np < 1$ then a graph will almost surely have no connected components with more than than $O(\log n)$ vertices,
- if $np = 1$ then a graph will almost surely have a largest component whose number of vertices is of order $n^{\frac{2}{3}}$
- if $np > 1$ then a graph will almost surely have a large connected component and no other component will contain more then $O(\log n)$ vertices.

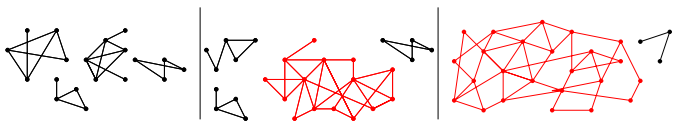

Fig. 6. The first regime models a graph when $np < 1$, middle when $np = 1$ and the last one when $np > 1$

The first regime was used because generating queries from smaller connected components theoretically makes intersections between them more probable. As it is expected that the queries have some underlying structure, the random graph model is an appropriate method of generating queries. The $np < 1$ case is helpful in this setting because the graph is naturally partitioned into several components as expected. The random graph could, for example, represent a database that contains data from many organizations, so each organization's data is a connected component and is separate from the others. `QueryGeneration` algorithm (Algorithm 3) is as follows.

Algorithm 3 Query Generation Algorithm

Require: N_Q, $Q = \emptyset$
for $|Q| \neq N_Q$ **do**
  $q = \emptyset$, $K = \emptyset$
  $l \leftarrow rand(minQueryLen, maxQueryLen)$
  $x = rand(v \in V)$
  $q \leftarrow q \cup x$
  $K \leftarrow K \cup \{v \mid (v, x) \in E, v \in V, x \in V\}$
  while $|q| \neq l$ **do**
    $x = rand(v \in K)$
    if $x \notin q$ **then**
      $q \leftarrow q \cup x$
    end if
  end while
  $Q \leftarrow Q \cup q$
end for

First, build a random graph G with $np < 1$. Key idea here is to generate random subgraph with number of nodes equal to l where l is the query length, such that $(minQueryLen \leq l \leq maxQueryLen)$. Repeat this till we generate desired number of queries. To assess the quality of our synthetic query workload generator, every query is compared with every other query to determine the size of intersections. Then the same number of queries were generated uniformly randomly and pairwise intersections were calculated. As expected, queries generated using `QueryGeneration` algorithm have much more intersections then queries generated randomly.

Real-world Dataset: We consider TREC Category B Section 1 dataset which consists of 50 million English pages. For queries we consider 40 million AOL search queries. In order to process these 50 million documents to document shards, we perform K-means clustering using Apache Mahout where K=10000. We consider each document shard as a data item in this paper. These document shards are distributed across 50 homogenous machines and are 3-way replicated. Each AOL web query is run through Apache Lucene to get top 20 document shards. Then we run our incremental set cover based routing to route queries to appropriate machines containing relevant document shards.

Overall, we evaluate our algorithms on a set of 50K synthetically generated queries generated from a graph with $np = .993$ and on real-world dataset. 20K queries from synthetic dataset and 8M queries from real-world dataset among them are used to create clusters and our routing approach is tested on remaining 30K queries from synthetic dataset and 32M queries from real-world dataset.

2) Baseline: When a query Q is received a request is sent to all machines that contains an element of Q. The machines are added to the set cover by the order in which they respond, until the query is covered. The first machine to respond is automatically added to the cover. The next machine to respond is added if it contains any element from the query that is not

yet in the cover. This process is continued until all elements of the query Q are covered. We call this method *baseline* set covering. While the method is fast, there is no discrimination in the machines taken, which means that the solution returned is far from optimal, as the next example illustrates:

Consider a query, $Q = \{1, \ldots n\}$, and a set of machines, $M_1, \ldots, M_{n+1}$, where $M_i = \{i\}$ for $i \leq n$ and $M_{n+1} = \{1, \ldots n\}$. If machines $M_1, \ldots, M_n$ respond first then this algorithm will cover Q with n machines where the optimal cover contains only one machine, M_{n+1}. Given N queries, our algorithm should improve upon the average optimality offered by this baseline covering method and be faster than running the greedy set cover N times.

3) Machine: The experiments were run on a Intel Core i7 quad core with hyperthreading CPU 2.93GHz, 16GB RAM, Linux Mint 13. We create multiple logical partitions within this setup and treat each logical partition as a separate machine. This particular experimental setup does not undermine the generality of our approach in anyway. Our results and observations stand valid even when run on distributed setup with real machines.

B. Experimental Analysis of Clusterings

We ran the the clustering algorithm on several sets of queries. All of these query sets are of size $50,000$ and are generated via the Erdős-Rényi graph regime with $0.9 < np < 1.0$ according to the query generation algorithm described in Section VII-A1. We specifically tested the resulting clusters for quality of clustering and for its applicability towards real-time processing. The ideas behind real-time processing are more thoroughly discussed in Section VI, but the essential idea is this: given a small sample of queries beforehand for pre-computing, we need to be able to process new queries as they arrive, with no prior information about them.

1) Clustering Quality: In a high-quality cluster most of the data elements have probability close to 1.0. Intuitively, this indicates that the queries in the cluster are all extremely similar to each other (i.e. they all contain nearly the same data elements). Then one measure of clustering quality would be to look at the probability of data elements across clusters.

As a first measure, we recorded the probability of each data element in each cluster, and Figure 8(a) depicts the results in a histogram for a typical clustering. The high frequency of data elements with probability over 0.9 indicates that a significant number of data elements have high probability within their cluster. (Note that in this analysis, if a data element is contained in many clusters, its probability is counted separately for each cluster.) However, interestingly, the distribution then becomes relatively uniform for all the other ranges of probabilities. This potentially illustrates the difference between mediocre and high-quality clusters described in Section IV-B. A more ideal clustering algorithm would increase the number of data elements in the higher probability bins and decrease the number of those in lower quality bins. Still, the prevalence of elements with probability greater than 0.9 is heartening because this indicates a fairly large common intersection among all the clusters. By processing this common intersection alone, we potentially cover a significant fraction of each query in the cluster with just a single greedy algorithm.

To paint a broad picture, the above measure of cluster quality ignores the clusters themselves. There may be variables inherent to the cluster which affect its quality and are overlooked. For example, perhaps clusters begin to deteriorate once they reach a certain size.

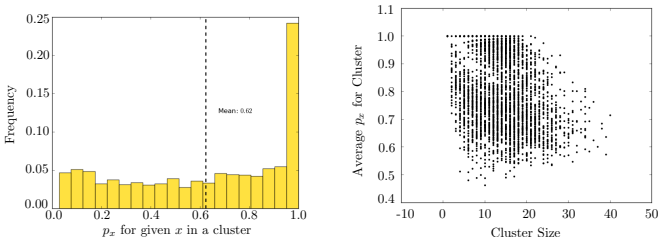

(a) Overall averages (b) Weighted by cluster

Fig. 8. These graphs are generated by clustering 50K synthetically generated queries for which $np = 0.973$. In Figure 8(a), we look at each data element in each cluster and record its probability in that cluster. A high-quality algorithm would have tall bars on the right and very short bars elsewhere. In Figure 8(b), we take a weighted average of the probability of each element. Note the downward trend as clusters get large. This may mean we should restrict cluster size.

Let us define the *average probability* of cluster K as:

$$\overline{p}(K) = \frac{1}{\sum_{Q \in K} |Q|} \sum_{Q \in K} \sum_{x \in Q} p_x \tag{9}$$

Essentially, $\overline{p}(K)$ is a weighted average of the probabilities of each element in the cluster (so data elements that are in many queries are weighted heavily). In Figure 8(b) we see that there is some deterioration of average probability as the clusters get larger, but for most of the size range, the quality is well scattered. While most of the clusters have average probability greater than 0.6, a stronger clustering algorithm would collapse this distribution upwards.

2) Real-time Applicability: For `simpleEntropy` to be successful in dealing with real-time queries we have additional requirements. First of all, the incoming queries need to be processed quickly. If, for example, the query needed to be checked against every single cluster before it was put in one, simply running the greedy algorithm on it would probably be faster, since there are on the order of thousands of clusters. The *fast* version of the algorithm, which only samples one element from each of the real-time queries, meets this requirement. Second, we want most of the cluster to be generated when only a small fraction of the queries are already processed. This way, most of the information about incoming queries is already computed, which allows us to improve running time. In Figure 9 and Table II, we see that more than 75% of the total clusters are generated with only 20% of the data processed, which means our clustering algorithm is working as we want it to. Finally, we want incoming queries to contain most of the high probability elements of the cluster. Specifically, let $Q_1, \ldots, Q_n$ be the queries in a cluster K, where $X = Q_1 \cap \cdots \cap Q_n$, and let Q^* be the incoming query. We want $X \subset Q^*$, since this means that the deepest G-part can cover a lot of Q with little waste. While this seems to be true for the high-quality clusters described above, we could seek to improve our algorithm to generate fewer mediocre clusters.

C. Experimental Comparison of Cluster Processing Algorithms

For comparing our cluster processing algorithms we had implemented two reference algorithms and two that we developed ourselves. We show in this section that our algorithms are successful, in that they are both fast and optimal.

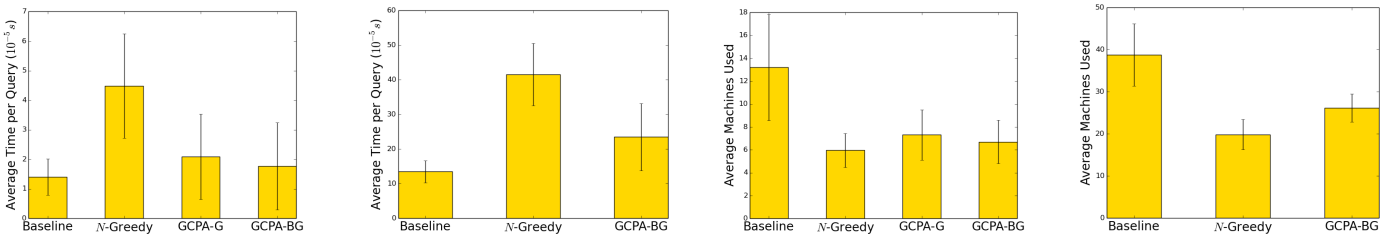

(a) Run-time for synthetic dataset (b) Run-time for real-world dataset (c) Optimality for synthetic dataset (d) Optimality for real-world dataset

Fig. 7. Comparison of run-time and optimality (average *query span*) of our algorithms on synthetic dataset and real-world dataset.

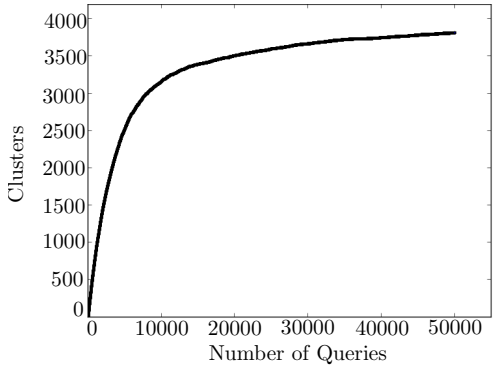

Fig. 9. We plot the number of clusters as each of 50k queries comes in, after being generated from an Erdős-Rényi graph with $np = .999$.

% Queries Processed	6.0	10.0	13.8	25.0	33.7	40.0	50.0	53.7	75.0	88.2	90.0	99.5
% Clusters Formed	50.0	66.1	75.0	86.6	90.0	91.9	94.3	95.0	97.9	99.0	99.2	99.9

TABLE II. THESE RESULTS CAN HELP SELECT THE THRESHOLD WE NEED FOR PRE-PROCESSING QUERIES TO FORM MOST OF OUR CLUSTERS. THE DATA ARE TAKEN FROM PROCESSING 50,000 QUERIES FROM A RANDOM GRAPH WITH $np = .999$. POTENTIAL THRESHOLDS INCLUDE $13.8\%, 33\%, 40\%$.

The first reference algorithm is the one primarily evaluated in the papers by Kumar and Quamar et al., [1][2][5], we call it `N_Greedy`. This is simply running the greedy algorithm on each query independently. This algorithm has the opposite properties of the baseline algorithms. While its covers are as close to optimal as possible, it has a longer run-time than the baseline. Thus, we want our algorithm to run faster than `N_Greedy`.

The two algorithms that we have developed and implemented are the `GCPA` with the greedy algorithm (`GCPA_G`) and `GCPA` with `BetterGreedy` (`GCPA_BG`). The major difference between the reference algorithms and our algorithms is that we are using clustering to exploit the correlations and similarities of the incoming queries. Our algorithms are faster than `N_Greedy` and more optimal than the baseline algorithms.

We compare the run-time and optimality (average number of machines that a query touches) of our algorithms and the two reference algorithms on both synthetic and real-world datasets. Our algorithms perform considerably faster than `N_Greedy` and are also faster than the smarter baseline algorithm. In terms of optimality, both our algorithms considerably outperform the standard baseline algorithm as shown in Figure 7. In summary, when evaluating with synthetic dataset, as shown in Figures 7(a) and 7(c), our technique is about $2.5\times$ faster when compared to repeated greedy technique `N_Greedy` and selects 50% fewer machines when compared to baseline routing technique. On the other hand, we evaluate on GCPA-BG for the real-world dataset because it has better optimality, and in the real-time case, the time penalty for using GCPA-BG over GCPA-G is only relevant in the pre-computing stage. For real-world dataset case, as shown in Figures 7(b) and 7(d), our technique is about $2\times$ faster when compared to repeated greedy technique `N_Greedy` and selects 32% fewer machines when compared to baseline routing technique. The error bars shown are one standard deviation long. Even though the figures show the results for only one set of queries, we have run dozens of samples, and the overall picture is the same. The results of our experiments provide strong indication that our algorithm is indeed an effective method for incremental set cover, in that it is faster than `N_Greedy` and more optimal than the baseline.

In terms of optimality, it is also important to do a pairwise comparison of cover lengths (i.e. does our algorithm perform better for queries of any size). Taking the average as we have done in Figure 7(c) masks potentially important variation. We want to ensure that our algorithm effectively handles queries of all sizes. In Figures 10(a) and 10(b), we compare the query-by-query performance (in terms of optimality) of our two algorithms against `N_Greedy` for the synthetic dataset. The x-axis is the number of sets required to cover a query using `N_Greedy`. The y-axis, "Δ Cover Length", is the length of the cover given by our algorithm minus the length of the greedy cover. The number next to the y-axis at $y = k$ shows the normalized proportion of queries for which the `GCPA` cover is at most k machines larger than the `N_Greedy`. The size of the circle indicates the number of queries at that coordinate. With `GCPA_BG`, we see that more than 90% of all queries are covered with at most one more machine than the greedy cover, and for the majority of queries, the covers are the same size. The `GCPA_G` algorithm does not perform quite as well. Even in this case, the majority of queries are covered using only one more machine than the greedy cover. Since `GCPA_BG` is slower than `GCPA_G`, users can choose their algorithm based on their preference for speed or optimality.

On the other hand, in Figure 10(c), we evaluate the performance of our real-time algorithm on the real-world dataset on a query-by-query basis. For each query, we record the number of machines used to cover it using our algorithm and the number of machines required to cover it using the baseline algorithm and record the difference, i.e. the "Δ Cover Length" on the y-axis is the size of the baseline cover minus the size of our algorithm's cover. The area of the circle at point (x, y) is proportional the number of queries for which our algorithm

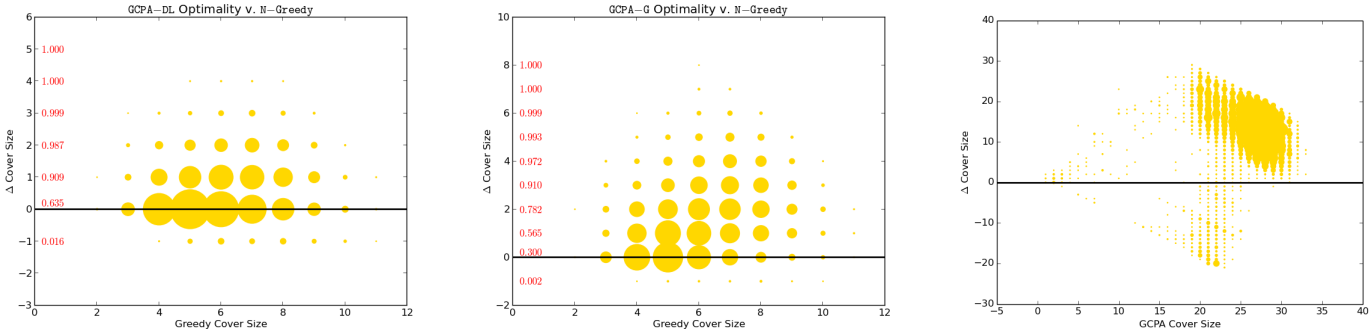

(a) `GCPA_DL` vs. `N_Greedy` for synthetic dataset (b) `GCPA_G` vs. `N_Greedy` for synthetic dataset (c) `GCPA` vs. `Baseline` for real-world dataset

Fig. 10. Pairwise comparisons of optimality for our algorithms.

used x machines to cover and for which the difference in cover length is y. Thus, the total area of the points above the $y = 0$ line represents where our algorithm outperforms the baseline algorithm. We see in the figure that the vast majority (96.5%) of the queries are covered more efficiently by our algorithm than by the baseline algorithm. This is actually a bit of an understatement because most of the queries for which the baseline algorithm performed with better optimality (i.e. points below $y = 0$) were queries of length one (or at least very small), which can easily be handled as special cases. In this case, only one machine is required to cover, but our algorithm takes the entire cover from the cluster that the length one query was put into. This situation is easily remedied. For small enough queries, especially queries of length one, instead of running our algorithm we should just cover them directly. In conclusion, we have delivered an algorithm that is significantly faster than `N_Greedy` and also more optimal than the baseline algorithm.

VIII. Conclusion

In this paper, we presented an efficient routing technique using the concept of *incremental set cover* computation. The key idea is to reuse the parts of *set cover* computations for previously processed queries to efficiently route real-time queries such that each query possibly touches a minimum number of machines for its execution. To enable the sharing of *set cover* computations across the queries, we take advantage of correlations between the queries by clustering the known queries and keeping track of computed query *set covers*. We then reuse the parts of already computed *set covers* to cover the remaining queries as they arrive in real-time. We evaluate our techniques using both real-world TREC with AOL datasets, and simulated workloads. Our experiments demonstrate that our approach can speedup the routing of queries significantly when compared to repeated greedy *set cover* approach without giving up on optimality. We believe that our work is extremely generic and can benefit variety of scale-out data architectures such as distributed databases, distributed IR, map-reduce, and routing of VMs on scale-out clusters.

References

[1] K. A. Kumar, A. Quamar, A. Deshpande, and S. Khuller, "Sword: Workload-aware data placement and replica selection for cloud data management systems," *The VLDB Journal*, vol. 23, no. 6, pp. 845–870, Dec. 2014.

[2] A. Quamar, K. A. Kumar, and A. Deshpande, "Sword: Scalable workload-aware data placement for transactional workloads," in *Proceedings of the 16th International Conference on Extending Database Technology*, ser. EDBT '13, 2013, pp. 430–441.

[3] C. Curino, E. Jones, Y. Zhang, and S. Madden, "Schism: A workload-driven approach to database replication and partitioning," *VLDB*, vol. 3, no. 1-2, pp. 48–57, Sep. 2010.

[4] A. Kulkarni and J. Callan, "Selective search: Efficient and effective search of large textual collections," *ACM Trans. Inf. Syst.*, vol. 33, no. 4, pp. 17:1–17:33, 2015.

[5] A. K. Kayyoor, "Minimization of resource consumption through workload consolidation in large-scale distributed data platforms," *Digital Repository at the University of Maryland*, 2014.

[6] V. V. Vazirani, *Approximation Algorithms*. Springer Science & Business Media, 2013.

[7] N. Alon, B. Awerbuch, Y. Azar, N. Buchbinder, and J. Naor, "The online set cover problem," *SIAM J. Comput.*, vol. 39, no. 2, pp. 361–370, 2009.

[8] A. Levin, "Approximating the unweighted k-set cover problem: Greedy meets local search," in *Approximation and Online Algorithms*, 2007, vol. 4368, pp. 290–301.

[9] C. C. Aggarwal and C. K. Reddy, *Data Clustering: Algorithms and Applications*. CRC Press, 2013.

[10] J. Kleinberg and E. Tardos, *Algorithm Design*. Addison-Wesley Longman Publishing Co., Inc., 2005.

[11] R. Baeza-Yates, C. Hurtado, and M. Mendoza, "Query recommendation using query logs in search engines," in *Current Trends in Database Technology - EDBT 2004 Workshops*, 2005, vol. 3268, pp. 588–596.

[12] S.-L. Chuang and L.-F. Chien, "Towards automatic generation of query taxonomy: a hierarchical query clustering approach," in *Data Mining, 2002. ICDM 2003. Proceedings. 2002 IEEE International Conference on*, 2002, pp. 75–82.

[13] "Query clustering using user logs," *ACM Trans. Inf. Syst.*, vol. 20, no. 1, pp. 59–81, 2002.

[14] V. T. Paschos, "A survey of approximately optimal solutions to some covering and packing problems," *ACM Computing Surveys (CSUR)*, vol. 29, no. 2, pp. 171–209, 1997.

[15] Z. Zhao, R. Song, X. Xie, X. He, and Y. Zhuang, "Mobile query recommendation via tensor function learning," in *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, 2015, pp. 4084–4090.

[16] N. Gupta, L. Kot, S. Roy, G. Bender, J. Gehrke, and C. Koch, "Entangled queries: enabling declarative data-driven coordination," in *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2011, Athens, Greece, June 12-16, 2011*, 2011, pp. 673–684.

[17] D. Barbará, Y. Li, and J. Couto, "COOLCAT: An entropy-based algorithm for categorical clustering," in *Proceedings of the eleventh international conference on information and knowledge management*. ACM, 2002, pp. 582–589.

[18] P. Erdős and A. Rényi, "On the evolution of random graphs," *Publications of the Mathematical Institute of the Hungarian Academy of Sciences*, vol. 5, pp. 17–61, 1960.